An Infinite Dimensional Model for A Single Server Priority Queue

Neal Master, Zhengyuan Zhou, and Nicholas Bambos

arXiv:1609.07996v1 [math.PR] 26 Sep 2016

Abstract— **We consider a Markovian single server queue in which customers are preemptively scheduled by exogenously assigned priority levels. The novelty in our model is that the priority levels are randomly assigned from a continuous probability measure rather than a discrete one. Because the priority levels are drawn from a continuum, the queue is modeled by a measure-valued stochastic process. We analyze the steady state behavior of this process and provide several results. We derive a measure that describes the average distribution of customer priority levels in the system; we provide a formula for the expected sojourn time of a customer as a function of his priority level; and we provide a formula for the expected waiting time of a customer as a function of his priority level. We interpret these quantitative results and give a qualitative understanding of how the priority levels affect individual customers as well as how they affect the system as a whole. The theoretical analysis is verified by simulation. We also discuss some directions of future work.**

I. Introduction

Priority queueing models are useful in a variety of different applications. In communication engineering, priority queues are used to study networks with differentiated levels of quality of service [1], [2]. In healthcare, priority queueing models are used to study and understand triage policies [3] in which certain types of patients are prioritized over others. In mathematical finance, limit orders are given priority for being matched with market orders according their price and time of arrival at the exchange [4]. Because of the breadth of the potential applications, many priority queueing models exist; see [5] for a standard reference on priority queueing models.

In this paper, we formulate and analyze a single server Markovian priority queueing model with the following novelty: we consider a *continuum* of priority levels. Because priority levels are uncountably infinite, unlikely previously studied models, e.g. [6], [7], our model requires an infinite dimensional state. As a result, standard Markov chain techniques that apply when there are finitely many priority levels, e.g. [8], do not apply. Although we restrict ourselves to a Markovian model with Poisson arrivals and exponential service times, the key difference is that our state is a function with an uncountable domain rather than merely being a finite dimensional vector.

Because of the complexity that arises due to this infinite dimensional state, we opt to simplify other aspects of the model. In particular, we assume that all customers experience the same service rate regardless of their priority level. This differs from previous work, e.g. [9], and restricts our attention to models in which differing priority levels affect the order in which jobs are scheduled but not the service rates that they experience. We also focus on the case of preemptive scheduling as in [10], [11] rather than non-preemptive scheduling as in [12], [13]. This simplifies our analysis because with preemptive scheduling we know that the customer who is being served is always the customer with the highest priority.

We note that the use of function-valued or measure-valued stochastic processes is itself not novel to queueing. Infinite dimensional models have been used to study the earliest-deadline-first discipline [14], the processor-sharing discipline [15], as well as many-server [16], [17] and infinite server [18] queueing models. In these contexts, the measure encodes dynamic properties of the jobs in the system such as their residual service times. In our model, because the priority levels are static, the state is constant between arrival and departure events. Consequently, our model is far more tractable. In fact, these other models focus on diffusion approximations while we present only exact results.

The measure-valued queueing model that is most closely related to our model can be found in [19]. This model was originally proposed to understand fragmentation is disk storage [20]. The model consists of an infinite server system in which servers (rather than customers) are ranked. This can be seen as servers with countably infinite priority levels. Our model is quite different because we study a single server system in which customers have continuous priority levels, but because both models have static priority levels, our styles of analysis are somewhat similar. The most significant divergence between our results is that we are able to provide exact results while [19] focuses on diffusion approximations.

With this background and motivation in mind, the remainder of the paper is organized as follows. In Section II we fully describe our model and discuss different choices for the state representation. In Section III we analyze the steady state behavior of the system. In particular, we derive a measure that tells us the average distribution of customer priority levels in the system. We also derive a formula for the average sojourn time of a customer as a function of his priority level. In Section IV we provide a simulation that verifies our analytical formulae. In Section V we discuss some potential directions of future work and we conclude in Section VI.

N. Master and Z. Zhou are each supported by a Stanford Graduate Fellowship (SGF) in Science & Engineering.

N. Master, Z. Zhou, and N. Bambos are with the Department of Electrical Engineering, Stanford University, Stanford, CA, 94305, USA. {nmaster, zyzhou, bambos}@stanford.edu

II. Model Formulation

In this section we formally describe our model. We explain our assumptions and highlight the fact that certain seemingly limiting assumptions are actually without loss of generality. We present three infinite dimensional state representations and explain their equivalence.

We consider a single server queue with an infinite buffer. Customers arrive according to a Poisson process with rate $\rho > 0$. We do not assume any upper bound on ρ. Customers have independent and identically distributed (IID) exponential service times. Time can be scaled arbitrarily so we assume that the service times have unit mean. In addition, customers have IID priority levels that are uniformly distributed on the unit interval. The priority levels are independent of all other random quantities in the model. Customers are scheduled preemptively according to their priorities: the highest priority customer will always be served even if this interrupts the service of another customer. An interrupted customer will wait in the queue until it is rescheduled for service, i.e. when it has the highest priority level of all the customers in the system. In summary, we have an $M/M/1$ queue (not necessarily stable) in which customers are preemptively scheduled according to exogenously assigned IID $U([0,1])$ priority levels.

Note that because the customers are scheduled based on their relative order rather than their absolute value, the fact that the priority levels are drawn from $U([0,1])$ (as opposed to some other distribution) is actually without loss of generality. Because the scheduling decisions only depend on the relative order of the priority levels, the dynamics would be unchanged if the priorities were transformed by any monotone map. In particular, suppose we want the priority levels to be drawn from some other distribution with cumulative distribution function (CDF) $F(\cdot)$. Consider customers i and j with priority levels p_i and p_j drawn from $U([0,1]$. Let $\tilde{p}_i = F^{-1}(p_i)$ and $\tilde{p}_j = F^{-1}(p_j)$ where $F^{-1}(\cdot)$ is the quantile function associated with $F(\cdot)$:

$$F^{-1}(p) = \inf\{x \in \mathbb{R} : p \leq F(x)\} \tag{1}$$

If $p_i > p_j$ then we will also have $\tilde{p}_i \geq \tilde{p}_j$. We also have that $\tilde{p}_i$ and $\tilde{p}_j$ are distributed according to the CDF $F(\cdot)$ [21, Theorem 2.1]. So if $F(\cdot)$ is strictly increasing then using $\tilde{p}_i$ and $\tilde{p}_j$ yields the same scheduling dynamics as using p_i and p_j. If $F(\cdot)$ is not strictly increasing, then with non-zero probability we could have $\tilde{p}_i = \tilde{p}_j$. However, in this situation customers i and j are indistinguishable and these ties can be broken in an arbitrary fashion, e.g. randomly. Consequently, our model encompasses arbitrary distributions of priority levels. For simplicity, we will focus having priority levels drawn from $U([0,1])$.

We also note that because of the memorylessness property of the exponential distribution, after a customer is preempted its residual service time is still exponentially distributed with unit mean. As a result, the state does not need to include the residual service time of each customer in the system, just the priority level of each customer. Since we have a continuum of priority levels, the state needs to encode the priority level of each customer in the system. It is convenient to encode this list of priority levels as a point measure on $[0,1]$. Let $\mathcal{B}([0,1])$ be the σ-algebra of Borel sets on $[0,1]$. Given $B \in \mathcal{B}([0,1])$ let $x_t(B)$ be the number of customers in the system at time t with priority levels contained in B. In other words, if there are N customers in the system at time t and their priority levels are $\{p_1, \ldots, p_N\} \subset [0,1]$, then

$$x_t = \sum_{i=1}^{N} \delta_{p_i} \tag{2}$$

where δ_z denotes a Dirac measure at $z \in [0,1]$.

We can equivalently represent the state by either the (non-normalized) CDF or the complementary CDF:

$$X_t(p) = x_t([0,p]), \quad \bar{X}_t(p) = x_t((p,1]) \tag{3}$$

The equivalence of these state representations follows from the fact that $\{[0,p] : p \in [0,1]\}$ and $\{(p,1] : p \in [0,1]\}$ each form π-systems that generate $\mathcal{B}([0,1])$. Because $x_t(\cdot)$ is a counting measure, we know that $x_t([0,1])$ is finite for all t. Hence, we can apply the π-λ Theorem to show that $\{X_t(p) : p \in [0,1]\}$ and $\{\bar{X}_t(p) : p \in [0,1]\}$ each uniquely define $x_t(\cdot)$. The definitions of π-systems and λ-systems along with the method of uniquely extending a measure from a π-system to a σ-algebra are standard in measure theory. For a reference, see [22, Chapter 3].

III. Some Theoretical Results

We now analyze the steady state behavior of the system. First we characterize the equilibrium distribution of $\bar{X}_t(p)$ for each $p \in [0,1]$. We provide a corollary that partially characterizes the equilibrium distribution of $x_t(\cdot)$. We then provide a formula for the expected sojourn time of a customer as a function of its priority level. As a small corollary to this we provide a formula for the expected waiting time of a customer as a function of its priority level. As in the previous section, we rely on standard results regarding the extension of measures from π-systems to σ-algebras which can be found in [22, Chapter 3].

Theorem 1: Fix any $p \in [0,1]$, $\bar{X}_t(p)$ converges weakly to a random variable $\bar{X}(p)$. If $(1-p)\rho < 1$, then $\bar{X}(p)$ is a geometrically distributed random variable on the non-negative integers with

$$\mathbb{E}[\bar{X}(p)] = \frac{(1-p)\rho}{1-(1-p)\rho}. \tag{4}$$

In other words, if $(1-p)\rho < 1$ then

$$\lim_{t\to\infty} \mathbb{P}(\bar{X}_t(p) = k) = (1-(1-p)\rho)((1-p)\rho)^k \tag{5}$$

for each non-negative integer k.

If $(1-p)\rho \geq 1$, then $\bar{X}(p) = \infty$ almost surely.

Proof: The key is to notice that because of the preemptive scheduling, the customers with priority levels in $(p,1]$ are not affected in any way by customers with lower

priority[1], i.e. the customers with priority levels in $[0, p]$. In addition, because the priority levels are independent of the inter-arrival times, the customers with priority levels in $(p, 1]$ arrive according to a Poisson process with rate $(1-p)\rho$. As a result, $\bar{X}_t(p)$ is stochastically equivalent to the population in an $M/M/1$ queue with unit service rate and arrival rate $(1-p)\rho$. As a result, $\bar{X}_t(p)$ converges weakly to a geometric random variable on the non-negative integers with the given mean [24, Chapter 3]. Because there is no upper bound on ρ, it is possible that $(1-p)\rho \geq 1$. In this case, the equivalent $M/M/1$ queue is not stable and hence $\bar{X}_t(p)$ diverges to infinity. ■

Corollary 1: Fix $B \in \mathcal{B}([0,1])$. Then $x_t(B)$ converges weakly to a random variable $x(B)$ with mean

$$\mu(B) = \mathbb{E}[x(B)] = \int_B m(p)dp \tag{6}$$

where $m(\cdot)$ is defined as follows:

$$m(p) = \begin{cases} \frac{\rho}{(1-(1-p)\rho)^2} & ,(1-p)\rho < 1 \\ \infty & ,(1-p)\rho \geq 1 \end{cases} \tag{7}$$

Proof: The previous theorem tells us that if $B = [a, b]$ for some $0 \leq a < b \leq 1$, then $x_t(B) = \bar{X}_t(a) - \bar{X}_t(b)$. Since $x_t([a, b])$ converges weakly to $\bar{X}(a) - \bar{X}(b)$, performing the integration gives us the same result as in the previous theorem. Indeed, note that for p such that $(1-p)\rho < 1$,

$$m(p) = -\frac{d}{dp}\left\{\frac{(1-p)\rho}{1-(1-p)\rho}\right\} = -\frac{d}{dp}\mathbb{E}[\bar{X}(p)] \tag{8}$$

Now note that intervals of this form are a π-system that generates $\mathcal{B}([0,1])$. Consequently, if $\rho < 1$ then $\mu([0,1]) < \infty$ and so this defines a unique measure on $\mathcal{B}([0,1])$. On the other hand, if $\rho \geq 1$, we can still extend the measure from the π-system to $\mathcal{B}([0,1])$, but uniqueness is no longer guaranteed. However, we can apply the same reasoning as above to define a unique measure on $\mathcal{B}([1-1/\rho, 1])$ where $\mu(\cdot)$ is finite. The fact that $\mu(B) = \infty$ for any B such that $B \cap [0, 1-1/\rho]$ has non-zero Lebesgue measure follows from the instability argument in the previous theorem. Hence, regardless of the value of ρ, we can conclude that the expression for the mean equilibrium behavior of $x_t(B)$ holds for any $B \in \mathcal{B}([0,1])$. ■

Remark 1: Although Theorem 1 gives the full distribution of $\bar{X}(p)$ for any p, Corollary 1 only characterizes the first-order statistics of $x(\cdot)$. This is because Theorem 1 does not characterize the joint distribution of $\{\bar{X}(p) : p \in [0,1]\}$. Because we only have the marginal distributions, i.e. the distribution of $\bar{X}(p)$ for a single p, the higher-order statistical behavior of $x(\cdot)$ does not follow from Theorem 1.

Because service can be preempted and hence customers can enter service multiple times, we formally define the terms "sojourn time" and "waiting time". In particular, we note that the amount of time a customer spends in service before being preempted is considered waiting.

[1]This same observation has been useful for other priority queueuing models [23].

Definition 1: A customer's sojourn time is the amount of time from when the customer arrives to when it departs after completing service.

Definition 2: A customer's waiting time is the amount of time from when the customer arrives to the beginning of the last time the customer enters service.

Theorem 2: Fix any $p \in [0,1]$ and let $s(p)$ be the expected sojourn time for a customer with priority p in steady state. Then if $(1-p)\rho < 1$ then

$$s(p) = \frac{1}{(1-(1-p)\rho)^2} \tag{9}$$

and if $(1-p)\rho \geq 1$ then $s(p) = \infty$.

Proof: Let $\bar{S}(p)$ be the average sojourn time for customers with priority levels in $(p, 1]$. Since the priority levels are uniformly distributed, the law of total probability tells us that

$$\bar{S}(p) = \int_p^1 s(q)\frac{1}{1-p}dq. \tag{10}$$

Because customers with priority levels in $(p, 1]$ arrive at a rate $(1-p)\rho$, Little's Law [25] tells us that

$$\mathbb{E}[\bar{X}(p)] = (1-p)\rho\bar{S}(p) = \rho\int_p^1 s(q)dq \tag{11}$$

which yields the following:

$$\int_p^1 s(q)dq = \begin{cases} \frac{1-p}{1-(1-p)\rho} & ,(1-p)\rho < 1 \\ \infty & ,(1-p)\rho \geq 1 \end{cases} \tag{12}$$

Differentiating gives us the result. ■

Corollary 2: Fix any $p \in [0,1]$ and let $w(p)$ be the expected waiting time for a customer with priority p in steady state to receive service. Then if $(1-p)\rho < 1$ then

$$w(p) = s(p) - 1 = \frac{1}{(1-(1-p)\rho)^2} - 1 \tag{13}$$

and if $(1-p)\rho \geq 1$ then $w(p) = \infty$.

Proof: The sojourn time is the sum of the waiting time and the service time. Since we have a unit service rate, we merely subtract 1 from $s(p)$ to get $w(p)$. ■

Remark 2: The functions $m(\cdot)$, $s(\cdot)$, and $w(\cdot)$ define key performance metrics for the system. As expected, each of these functions is decreasing: higher priority customers wait less and see a smaller backlog than lower priority customers. Moreover, if we ignore some constant factors, each the functions decays as $1/p^2$. Hence, we see that the benefits of having higher priority grow quadratically. For instance, consider a customer with priority level 1.0 and a customer with priority level 0.5. The higher priority customer only has twice the priority of the lower priority customer but the higher priorty customer has an expected sojourn time that is roughly a quarter of the lower priority customer's expected sojourn time.

Remark 3: If $\rho \geq 1$ then all of these results exhibit a bifurcation, i.e. a qualititative change in behavior, at

$$p^* = 1 - \frac{1}{\rho}. \tag{14}$$

It is intuitive that when the server is overloaded, lower priority customers will be ignored so that higher priority customers can be served. The quantity p^* makes this intuition precise: when the queue is overloaded, customers with priority levels in $[0, p^*]$ will have infinite expected waiting times while customers in $(p^*, 1]$ will have finite expected waiting times.

Remark 4: Because of the aforementioned birfurcation, the case of $\rho = 1$ is particularly interesting. We know that when $\rho = 1$ the $M/M/1$ is unstable. However, since $p^* = 0$, all customers with priority levels in $(0, 1]$ have a finite sojourn time and only customers with priority levels equal to zero have infinite sojourn times. This seems a bit paradoxical: the queue is unstable but almost every customer has a finite sojourn time. This counterintuitive result arises because $\rho = 1$ is the critical point between a stable $M/M/1$ queue and an unstable $M/M/1$ queue.

IV. Simulation Verification

In this section, we report the results of two discrete event simulations of the system: one with $\rho < 1$ and one with $\rho \geq 1$. In both cases, we simulate for a time horizon of $T = 10^4$ and use the simulated data to estimate $m(\cdot)$, $s(\cdot)$, and $w(\cdot)$. In general, we see that the estimates match our theoretical results, thus supporting our analysis.

A. Estimation Methods

We first outline our estimation methods. For each of the functions that we estimate, we take a non-parametric approach: first we get local estimates and we then linearly interpolate to estimate the entire function. The details for each function are outlined below.

We compute our estimate of $m(\cdot)$, which we denote $\hat{m}(\cdot)$, as follows:

1) Because of the PASTA property [26], we record $x_t(\cdot)$ as observed immediately before a new arrival.
2) For $p_i \in \{0.025 + 0.05i\}_{i=0}^{19}$, we average the number of customers with priority levels in the half-open interval $[p_i - 0.025, p_i + 0.025)$ across our observations and scale this average by $1/0.05$. This gives us $\hat{m}(p_i)$.
3) We linearly interpolate $\{m(p_i)\}_{i=0}^{19}$ to get a complete estimate of $m(\cdot)$.

We compute our estimate of $s(\cdot)$, which we denote $\hat{s}(\cdot)$, in a similar fashion:

1) We record the arrival time, the departure time, and the priority level of each customer. If a customer does not depart, then his departure time is infinite.
2) For $p_i \in \{0.025 + 0.05i\}_{i=0}^{19}$, we average the sojourn times for customers with priority levels in the half-open interval $[p_i - 0.025, p_i + 0.025)$. This gives us $\hat{s}(p_i)$.
3) We linearly interpolate $\{s(p_i)\}_{i=0}^{19}$ to get a complete estimate of $s(\cdot)$.

We compute our estimate of $w(\cdot)$, which we denote $\hat{w}(\cdot)$, in a similar fashion:

1) We record the arrival time and the priority level of each customer. We also record the last time that the customer enters service before departing. If the customer never departs them this time is infinite.
2) For $p_i \in \{0.025 + 0.05i\}_{i=0}^{19}$, we average the waiting times for customers with priority levels in the half-open interval $[p_i - 0.025, p_i + 0.025)$. This gives us $\hat{w}(p_i)$.
3) We linearly interpolate $\{w(p_i)\}_{i=0}^{19}$ to get a complete estimate of $w(\cdot)$.

B. Estimation Results

First consider a simulation for which $\rho = 0.75$. In this case, the queue is stable and so $m(\cdot)$, $s(\cdot)$, and $w(\cdot)$ are finite. The results are plotted in Fig. 1. We see that for $m(\cdot)$, $s(\cdot)$, and $w(\cdot)$, the estimates agree with our theoretical analysis. Moreover, if we see that $\hat{s}(\cdot)$ and $\hat{w}(\cdot)$ have the same shape and merely differ by a constant. This confirms our previous analysis regarding the mean equilibrium behavior of $x_t(\cdot)$, the expected sojourn time, and the expected waiting time.

Now consider a simulation for which $\rho = 1.25$. In this case, the queue is not stable and so $m(\cdot)$, $s(\cdot)$, and $w(\cdot)$ are finite only for $p \in (p^*, 1] = (0.2, 1]$. As a result, we do not plot the functions for $p < p^*$. Because of the vertical asymptote at p^*, we use a log-scale for the vertical axis. The results are plotted in Fig. 2. In Fig. 2a, we see that $\hat{m}(\cdot)$ and $m(\cdot)$ seem to agree on $(p^*, 1]$. Fig. 2a also depicts the bifurcation at p^*. We see that for $p \in \{p_i\}_{i=0}^{3}$, $\hat{m}(p)$ is roughly 10 times the value of $\hat{m}(p_4)$. This is because $p_3 < p^*$ while $p_4 > p^*$. We see similar results regarding $\hat{s}(\cdot)$ in Fig. 2b. For $p \in (p^*, 1]$, $\hat{s}(p)$ and $s(p)$ agree. Note that for $p < p^*$, neither $\hat{s}(p)$ nor $s(p)$ appear on the plot. This is because both quantities are infinite. Hence, we see that $\hat{s}(\cdot)$ and $s(\cdot)$ agree for all $p \in [0, 1]$. We see the same results for $\hat{w}(\cdot)$: the estimate agrees with the analytic result where both are finite and also where both are infinite.

V. Future Work

There are several potential avenues of future work. One is derive more results about this particular model. In particular, deriving results regarding the higher order statistics of $x(\cdot)$ would be interesting but it is not immediately clear how to do this given the present results.

It would also be interesting to extend this model to networks of queues. With a single queue, the state is a point measure on $[0, 1]$ but if there are n queues connected in a network, the state would be a point measure on $[0, 1]^n$. It seems likely that the steady state would have a product-form similar to Jackson's Theorem [27], but the details of the analysis are not immediately clear.

Another direction of future modeling work would be to consider how continuous priority levels affect queues with many servers. The style of analysis would be similar and it seems reasonable that there are analogous results that can be be derived.

Finally, we mention that heavy traffic analysis may yield some interesting results. Priority queues are a canonical example of a system that exhibits "state-space collapse" in heavy traffic [28]. Indeed, if upon appropriate rescaling, we

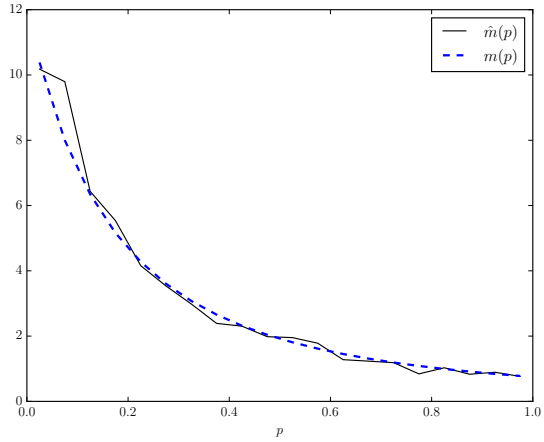

(a) $\hat{m}(\cdot)$ vs. $m(\cdot)$

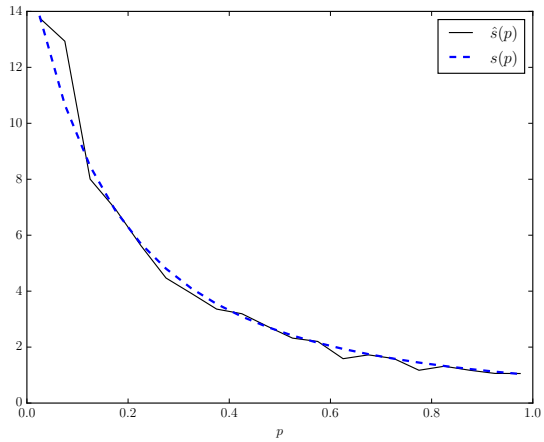

(b) $\hat{s}(\cdot)$ vs. $s(\cdot)$

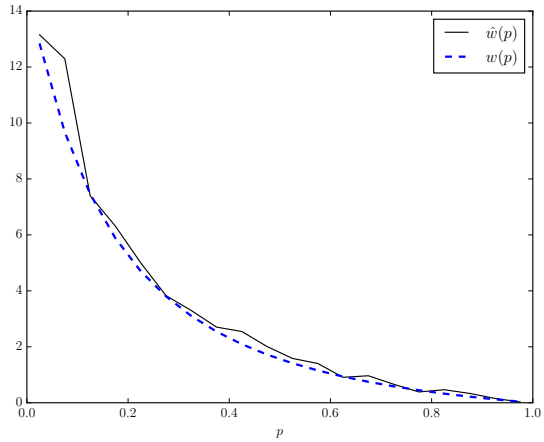

(c) $\hat{w}(\cdot)$ vs. $w(\cdot)$

Fig. 1: Estimates of $m(\cdot)$, $s(\cdot)$, and $w(\cdot)$ based on the data generated by simulating the system with $\rho = 0.75$ for a horizon of $T = 10^4$ time units. For this value of ρ, the queue is stable and so we use a linear scale for both axes.

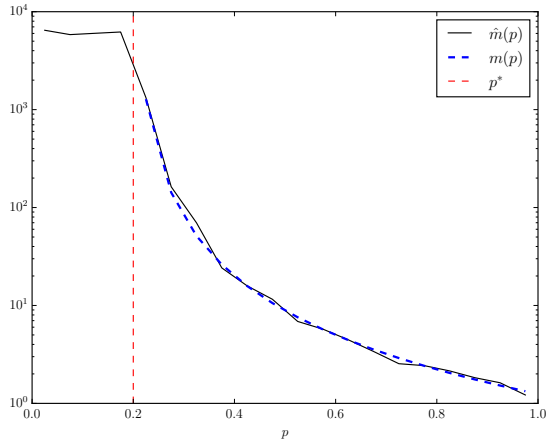

(a) $\hat{m}(\cdot)$ vs. $m(\cdot)$

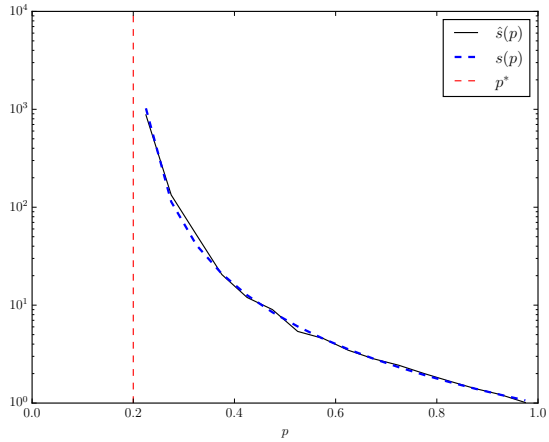

(b) $\hat{s}(\cdot)$ vs. $s(\cdot)$

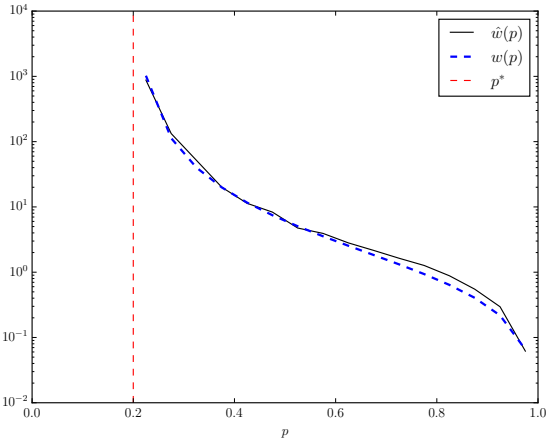

(c) $\hat{w}(\cdot)$ vs. $w(\cdot)$

Fig. 2: Estimates of $m(\cdot)$, $s(\cdot)$, and $w(\cdot)$ based on the data generated by simulating the system with $\rho = 1.25$ for a horizon of $T = 10^4$ time units. Because the queue is unstable, the values of the functions become quite large and hence we opt to use a logarithmic vertical axis.

would see that the diffusion limit associated with $X_t(p)$ for $p < p^*$ would be zero. However, in may be possible to take $\rho \uparrow \infty$ in such a way that $p^* \downarrow 0$ and we have are left with a non-trivial measure-valued diffusion. This idea is not yet well developed but because our analysis applied to overloaded queues, a heavy traffic limit is a natural idea to consider.

VI. Conclusions

We have presented an infinite dimensional model for a single server priority queue in which customers' priority levels are drawn from a continuous probability distribution. Our steady state analysis characterizes the mean behavior of the measure-valued process that describes the priority levels of the customers in the queue. We have also provided formulae for the expected sojourn and waiting times of customers as function of their priority levels. When the queue is overloaded, all of these analytical results exibit a bifurcation around a critical priority level. In particular, customers with priority levels strictly larger than this critical level will have finite expected sojourn times while customers with priority levels less than or equal to this critical level will have infinite expected sojourn times. To further bolster this analysis, we have also presented some simulations that agree with our formulae. We have also discussed some directions of future work.

References

[1] J. Shin, J. W. Kim, and C.-C. Kuo, "Quality-of-service mapping mechanism for packet video in differentiated services network," *IEEE Transactions on Multimedia*, vol. 3, no. 2, pp. 219–231, 2001.

[2] C. Semeria, "Supporting differentiated service classes: queue scheduling disciplines," *Juniper networks*, pp. 11–14, 2001.

[3] L. Green, "Queueing analysis in healthcare," in *Patient flow: reducing delay in healthcare delivery*, pp. 281–307, Springer, 2006.

[4] R. Cont, S. Stoikov, and R. Talreja, "A stochastic model for order book dynamics," *Operations research*, vol. 58, no. 3, pp. 549–563, 2010.

[5] N. K. Jaiswal, *Priority queues*, vol. 50. JSTOR, 1968.

[6] D. R. Miller, "Computation of steady-state probabilities for M/M/1 priority queues," *Operations Research*, vol. 29, no. 5, pp. 945–958, 1981.

[7] W. Whitt, "Weak convergence theorems for priority queues: preemptive-resume discipline," *Journal of Applied Probability*, vol. 8, no. 1, pp. 74–94, 1971.

[8] H. White and L. S. Christie, "Queuing with preemptive priorities or with breakdown," *Operations research*, vol. 6, no. 1, pp. 79–95, 1958.

[9] L. Takács, "Priority queues," *Operations Research*, vol. 12, no. 1, pp. 63–74, 1964.

[10] G. Yeo, "Preemptive priority queues," *Journal of the Australian mathematical society*, vol. 3, no. 04, pp. 491–502, 1963.

[11] W. Chang, "Preemptive priority queues," *Operations research*, vol. 13, no. 5, pp. 820–827, 1965.

[12] U. Bagchi and R. S. Sullivan, "Dynamic, non-preemptive priority queues with general, linearly increasing priority function," *Operations research*, vol. 33, no. 6, pp. 1278–1298, 1985.

[13] A. S. Kapadia, M. F. Kazmi, and A. C. Mitchell, "Analysis of a finite capacity non-preemptive priority queue," *Computers & operations research*, vol. 11, no. 3, pp. 337–343, 1984.

[14] B. Doytchinov, J. Lehoczky, and S. Shreve, "Real-time queues in heavy traffic with earliest-deadline-first queue discipline," *Annals of Applied Probability*, pp. 332–378, 2001.

[15] H. C. Gromoll, "Diffusion approximation for a processor sharing queue in heavy traffic," *Annals of Applied Probability*, pp. 555–611, 2004.

[16] G. Pang and W. Whitt, "Two-parameter heavy-traffic limits for infinite-server queues," *Queueing Systems*, vol. 65, no. 4, pp. 325–364, 2010.

[17] H. Kaspi and K. Ramanan, "SPDE limits of many-server queues," *The Annals of Applied Probability*, vol. 23, no. 1, pp. 145–229, 2013.

[18] J. Reed and R. Talreja, "Distribution-valued heavy-traffic limits for the $G/GI/\infty$ queue," *The Annals of Applied Probability*, vol. 25, no. 3, pp. 1420–1474, 2015.

[19] D. Aldous, "Some interesting processes arising as heavy traffic limits in an $M/M/\infty$ storage process," *Stochastic processes and their applications*, vol. 22, no. 2, pp. 291–313, 1986.

[20] E. G. Coffman, Jr, T. Kadota, and L. A. Shepp, "A stochastic model of fragmentation in dynamic storage allocation," *SIAM Journal on Computing*, vol. 14, no. 2, pp. 416–425, 1985.

[21] L. Devroye, "Sample-based non-uniform random variate generation," in *Proceedings of the 18th conference on Winter simulation*, pp. 260–265, ACM, 1986.

[22] P. Billingsley, *Probability and Measure*. John Wiley & Sons, 2008.

[23] P. J. Burke, "Priority Traffic with at Most One Queuing Class," *Operations Research*, vol. 10, no. 4, pp. 567–569, 1962.

[24] L. Kleinrock, "Queueing systems, Volume I: Theory," 1975.

[25] J. D. Little, "A proof for the queuing formula: $L = \lambda W$," *Operations research*, vol. 9, no. 3, pp. 383–387, 1961.

[26] R. W. Wolff, "Poisson arrivals see time averages," *Operations Research*, vol. 30, no. 2, pp. 223–231, 1982.

[27] J. R. Jackson, "Jobshop-like queueing systems," *Management science*, vol. 10, no. 1, pp. 131–142, 1963.

[28] M. I. Reiman, "Some diffusion approximations with state space collapse," in *Modeling and performance evaluation methodology*, pp. 207–240, Springer, 1984.